



IP Multicast Engineering for Access Grid Nodes

State of the Art in IPv4 Multicast
Deployment

Bill Nickless
nickless@mcs.anl.gov
<http://www.mcs.anl.gov/home/nickless>

Presented by: Bill Nickless



Problems Solved

- Distance Vector Metric Routing Protocol (DVMRP) does not scale
 - Easy to create IP Multicast “amplifiers”.
 - Separate tunneled routing infrastructure, not aligned with modern BGP Internetworking.
- Flood & Prune does not scale
 - Examples: PIM-Dense Mode, DVMRP.
 - Not sensitive to available bandwidth.
 - Requires receivers that are smart and powerful enough to send prune messages.



Pieces of the Solution

- Multicast Border Gateway Protocol.
 - Provides reachability and policy control for multicast routing, just as BGP does for unicast.
- Protocol Independent Multicast (Sparse Mode)
 - Listeners receive traffic only when requested.
 - Forms multicast distribution trees.
- Multicast Source Discovery Protocol
 - Finding active sources in other PIM Sparse Mode domains (usually other ASes).



Setting Policy: Multicast Border Gateway Protocol

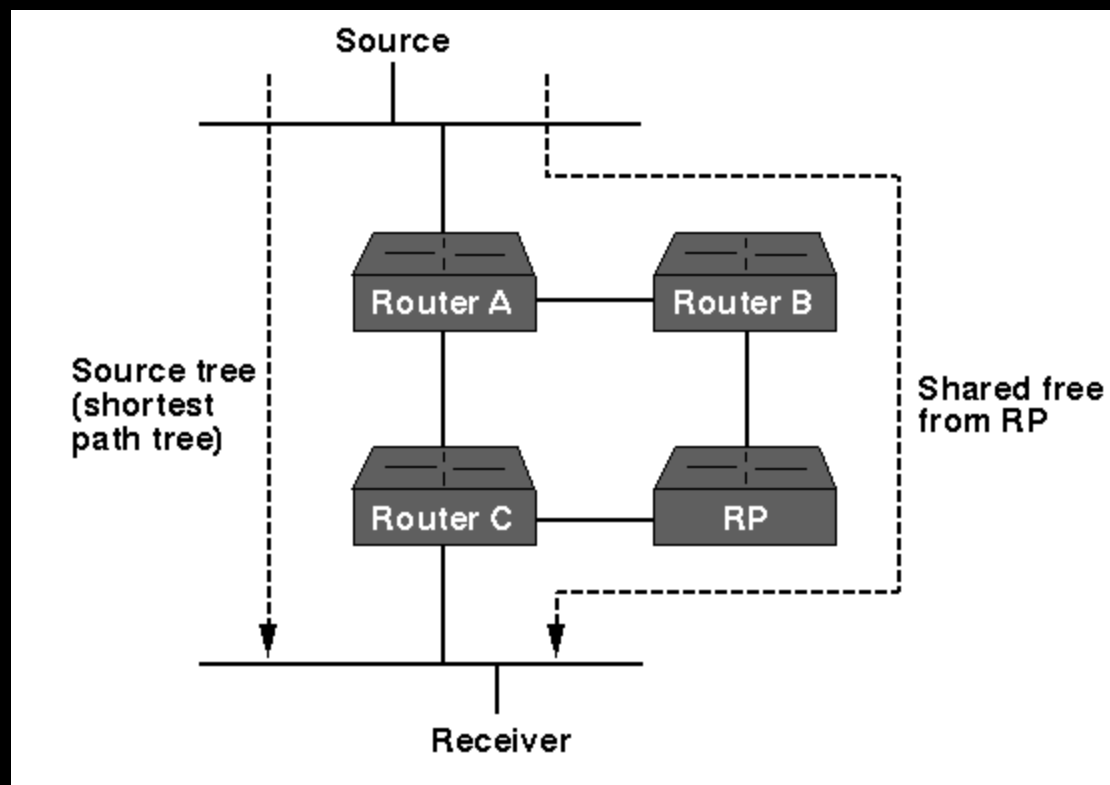
- RFC 2283 adds the MP_REACH_NLRI attribute to BGP-4.
 - Identifies a BGP route as unicast, multicast, or both unicast and multicast.
- When implemented in a router, all the standard BGP machinery is available for prefix filtering, preference setting, MEDs, AS length comparisons, etc.
- M-BGP routes can be independent of BGP, allowing for diverse inter-AS routing tables and policy.



PIM Sparse Mode

- RFC 2362 defines PIM Sparse Mode.
- No PIM-SM activity until a host starts transmitting traffic, or a host subscribes to a group.
- A Rendezvous Point (RP) is elected to become the root of the distribution tree for an IP Multicast Group within an AS.
- Given enough traffic, a source-based distribution tree is created.

PIM Sparse Mode





Multicast Session Discovery Protocol

- Not yet an RFC. See <http://www.ietf.org/html.charters/msdp-charter.html> and <http://www.ietf.org/internet-drafts/draft-ietf-msdp-spec-06.txt>
- PIM-SM RPs in separate ASes communicate through MSDP to find active multicast sources.
- Currently only covers IPv4. IETF's working group may decide to extend to IPv6.



The Beacon: Test Signal

- Testing Multicast requires active sessions
- <http://dast.nlanr.net/projects/beacon>
- In Java, so runs anywhere

Time: Mon Jan 22 17:38:10 CST 2001

Target: 233.2.171.1:56464

Beacons: 46 [details](#)

Page: refresh in 60 seconds

	Loss (%)	S0	S1	S2	S3	S4	S5	S6
R0	beacon@203.255.248.249	0	NA	0	NA	NA	0	0
R1	beacon@babar.switch.ch	NA	0	0	NA	NA	NA	NA
R2	beacon@hendrix.multicasttech.com	NA	NA	0	NA	25	40	19
R3	beacon@agaudio.arsc.edu	NA	NA	NA	0	NA	NA	NA
R4	beacon@agaudio.bu.edu	NA	NA	15	NA	0	0	0
R5	beacon@pgp1.cit.cornell.edu	NA	NA	22	NA	0	0	27
R6	beacon@palpatine.ucs.indiana.edu	NA	NA	2	NA	0	0	0



The Beacon: Issues

- Shows current state only.
 - Archive state over time?
 - How to visualize evolving state? Inherently a 3-dimensional problem, since state is 2D already.
- Server scaling problems with $O(40)$ beacons.
- Assumes Internet Standard Multicast model.



North American R&E Network IP Multicast Status

- ESNNet, Abilene, vBNS+, and NREN all running M-BGP, MSDP, and PIM-SM amongst themselves and with their customers/peers.
- Regional and Institutional networks are currently the most common limitations.
- STARTAP router in Chicago is an international IP multicast meeting point.
- International networks are coming online.